



Subscribe Register
(Full Service) (Limited Service, Free)

Login

Search: ☒ The ACM Digital Library ☐ The Guide

address <sentence> correlat* <sentence> cache

THE ACM DIGITAL LIBRARY

Enter words, phrases or names below. Surround phrases or full names with double quotation marks.

Search within Results: 18,632 found

address <sentence> correlat* <sentence> cache

□□□□

Desired Results:

must have **all** of the words or phrases

must have **any** of the words or phrases

must have **none** of the words or phrases

Only search in:*

☐ Title ☐ Abstract ☐ Review ☒ All Information

*Searches will be performed on all available information, including full text where available, unless spec

ISBN / ISSN: ☒ Exact ☐ Expand

Published:

By: ☒ all ☐ any ☐ none

In: ☒ all ☐ any ☐ none

Since:

Month



Year



Before:

December ▼

2000 ▼

As: Any type of publication ▼

Classification: (CCS) ☐ **Primary Only**Classified as: ☒ all ☐ any ☐ noneSubject Descriptor: ☒ all ☐ any ☐ noneKeyword Assigned: ☒ all ☐ any ☐ none

The ACM Portal is published by the Association for Computing

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#)



Subscribe Register
(Full Service) (Limited Service, Free)

Login

Search: ☒ The ACM Digital Library ☐ The Guide

THE ACM DIGITAL LIBRARY

Feedback

Published before December 2000

Sort results
by

relevance

☒ Save results to a Binder

Try

☒ Search Tips

Try

☐ Open results in a new window

Display results

expanded form

Results 1 - 11 of 11

1 Call graph prefetching for database applications

Murali Annavaram, Jignesh M. Patel, Edward S. Davidson

November 2003

ACM Transactions on Computer Systems (TOCS), Volume 2:

Full text available: pdf(701.71 KB)

Additional Information: full citation, abstract, reference

With the continuing technological trend of ever cheaper and larger memory, n soon be able to reside in main memory. In this configuration, the performance between the processing speed of the CPU and the memory access latency. Pre applications have large instruction and data footprints and hence do not use p paper, we propose Call Graph Prefetching (CGP), ...

Keywords: Instruction cache prefetching, call graph, database

2 Session 3: Energy-aware OS's: The benefits of event: driven energy accou

Frank Bellosa

September 2000

Proceedings of the 9th workshop on ACM SIGOPS European works
the operating system



Full text available: pdf(86.80 KB)

Additional Information: full citation, abstract, refere

A prerequisite of energy-aware scheduling is precise knowledge of any activity Embedded hardware monitors (e.g., processor performance counters) have pr field of performance analysis. The same approach can be applied to investigat individual threads. We use information about active hardware units (e.g., inte interface) gathered by event counters to establish a t ...

3 Energy efficient architectures: Reducing power requirements of instruction allocation of multiple datapath resources

Dmitry Ponomarev, Gurhan Kucuk, Kanad Ghose

Full text available:  pdf(1.51 MB)  Publisher Site

Additional Information: full citation, .

The "one-size-fits-all" philosophy used for permanently allocating datapath resources to maximize performance across a wide range of applications results in the over-allocation of resources. To reduce power dissipation in the datapath, the resource allocations can be dynamically adjusted based on the requirements of applications. We propose a mechanism to dynamically, simultaneously and efficiently allocate the instruction queue (IQ), the reorder buffer (R ...



Keywords: dynamic instruction scheduling, energy-efficient datapath, power r

4 Using a user-level memory thread for correlation prefetching

Yan Solihin, Jaejin Lee, Josep Torrellas

May 2002

ACM SIGARCH Computer Architecture News, Volume 30 Issue 2

Full text available:  pdf(1.48 MB)  Publisher Site

Additional Information: full citation, abstract

This paper introduces the idea of using a User-Level Memory Thread (ULMT) for correlation prefetching. In this approach, a user thread runs on a general-purpose processor in main memory in a DRAM chip. The thread performs correlation prefetching in software, sending requests to the cache of the main processor. This approach requires minimal hardware beyond what is already available. A software data structure that resides in main memory ...

Keywords: data prefetching, intelligent memory, processing-in-memory, compiler-aided prefetching, memory hierarchies, caches, threads

5 Memory-wall: Execution history guided instruction prefetching

Yi Zhang, Steve Haga, Rajeev Barua

June 2002

Proceedings of the 16th international conference on Supercomputing

Full text available:  pdf(218.17 KB)

Additional Information: full citation, abstract, references

The increasing gap in performance between processors and main memory has made prefetching techniques more important than ever. A major deficiency of existing prefetching techniques is the lack of an extra port to I-cache. A recent study by [19] shows that this factor alone explains why most microprocessors do not use such hardware-based I-cache prefetch schemes. To address this problem, we first present a method that does not require an extra port to I-cache ...

Keywords: hardware, instruction cache, performance, prefetching

6 Delayed Internet routing convergence

Craig Labovitz, Abha Ahuja, Abhijit Bose, Farnam Jahanian

June 2001 IEEE/ACM Transactions on Networking (TON), Volume 9 Issue 3

Full text available:  pdf(220.26 KB)

Additional Information: full citation, abstract, references,



This paper examines the latency in Internet path failure, failover, and repair of interdomain routing. Unlike circuit-switched paths which exhibit failover on the order of milliseconds, our experimental measurements show that interdomain routers in the packet-switched network take tens of minutes to reach a consistent view of the network topology after a fault. These delays are caused by fluctuations in the operation of the Border Gateway Protocol (BGP) ...

Keywords: Internet, failure analysis, network reliability, routing

7 Dead-block prediction & dead-block correlating prefetchers

An-Chow Lai, Cem Fide, Babak Falsafi

May 2001 ACM SIGARCH Computer Architecture News , Proceedings of the 28th Annual Symposium on Computer Architecture, Volume 29 Issue 2

Full text available:  pdf(972.60 KB)  Publisher Site

Additional Information: full citation, abstract, references,

Effective data prefetching requires accurate mechanisms to predict which cache blocks to prefetch and when to prefetch them. Dead-Block Predictors (DBPs), trace-based predictors that accurately predict which cache blocks will become evictable, enable placing data directly into L1, obviating the need for L2 prefetching.

8 Delayed Internet routing convergence

Craig Labovitz, Abha Ahuja, Abhijit Bose, Farnam Jahanian

August 2000 ACM SIGCOMM Computer Communication Review , Proceedings of the SIGCOMM 2000 Conference on Technologies, Architectures, and Protocols for Computer Communication

Full text available:  pdf(313.83 KB)


Additional Information: full citation, abstract, references,

This paper examines the latency in Internet path failure, failover and repair of inter-domain routing. Unlike switches in the public telephony network which exhibit failover on the order of milliseconds, our experimental measurements show that inter-domain routers take tens of minutes to reach a consistent view of the network topology after a fault. These delays are caused by temporary routing table oscillations formed during the operation of the Border Gateway Protocol (BGP) ...

9 A flow-based approach to datagram security

Suvo Mittra, Thomas Y. C. Woo

October 1997 ACM SIGCOMM Computer Communication Review , Proceedings of the ACM SIGCOMM Computer Communication Review , Applications, technologies, architectures, and protocols for computer communication

Full text available:  pdf(2.04 MB)


Additional Information: full citation, abstract, references, citation metrics

Datagram services provide a simple, flexible, robust, and scalable communication model. This has been well demonstrated by the success of IP, UDP, and RPC. Yet, the overwhelming majority of protocols that have been proposed are geared towards connection-oriented communication. Datagram communications tend to either rely on long term host-pair keying or require connection setup semantics. Separately, the

10 Scheduling and page migration for multiprocessor computer servers

Rohit Chandra, Scott Devine, Ben Verghese, Anoop Gupta, Mendel Rosenblum

November 1994 Proceedings of the sixth international conference on Architectural support for operating systems, Volume 29 , 28 Issue 11 , 5

Full text available:  pdf(1.56 MB)

Additional Information: full citation, abstract, references, citation metrics

Several cache-coherent shared-memory multiprocessors have been developed to support multiprogramming and parallel application workloads. Process scheduling and page migration policies on the performance of these machines. This paper describes scheduling and page migration policies on the performance of these machines.

11 Techniques for compressing program address traces

Andrew R. Pleszkun

November 1994 Proceedings of the 27th annual international symposium on Memory management

Full text available:  pdf(931.63 KB)

Additional Information: full citation, abstract, references, citation metrics




In this paper a technique for generating consistent, reproducible traces with a compression ratio better than standard general-purpose compression programs is described. An intermediate form is generated and then read as the input to the second program source code is required, and this technique will work on address streams. The way the address trace is encoded.

Keywords: compression, trace generation

Results 1 - 11 of 11

The ACM Portal is published by the Association for Computing Machinery. C

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

 Useful downloads:  Adobe Acrobat  QuickTime  Windows Media